U.S. DEPARTMENT OF ENERGY

OFFICE OF SCIENCE

Advanced Scientific Computing Research

Positioning for Exascale

Salishan
Conference
April 27, 2010

Walter M. Polansky
Advanced Scientific Computing Research
Office of Science, U.S. Department of Energy

ASCR Vision

Maintain position as "Best in Class" in advancing science and technology through modeling and simulation through

- **Continued excellence in applied mathematics and computer science research**
- **Strengthening new and established modes for cross-disciplinary partnerships**
- **Delivering Petascale Computing for Science Applications**
- **Realize the promise of Exascale**

ASCR Priorities

- To develop mathematical descriptions, models, methods and algorithms to enable scientists to accurately describe and understand the behavior of complex systems involving processes that span vastly different time and/or length scales
- To develop the underlying understanding and software to make effective use of computers at extreme scales
- To transform extreme scale data from experiments and simulations into scientific insight.
- To advance key areas of computational science and discovery that advance the missions of the Office of Science through mutually beneficial partnerships.
- To deliver the forefront computational and networking capabilities, enabling world-class researchers to extend the frontiers of science.
- To develop networking and collaboration tools and facilities that enable scientists worldwide to work together.

FY 2011 President's Request

ASCR Budget Summary	FY 2009 Appropriation	FY 2010 Appropriation	FY 2011 President's Request	Delta FY10 to FY11 Req.
Advanced Scientific Computing Research				
Applied Mathematics	45,161	44,792	45,450	+ 658
Computer Science	30,782	46,800	47,400	+ 600
Computational Partnerships (includes SciDAC)	59,698	53,293	53,297	+ 4
Next Generation Networking for Science	14,732	14,321	14,321	--
SBIR/STTR	—	4,586	4,623	+ 37
Total, Mathematical, Computational, and Computer Sciences Research	*150,373*	*163,792*	*165,091*	*+1,299*
High Performance Production Computing (NERSC)	53,497	55,000	56,000	+ 1,000
Leadership Computing Facilities	116,222	123,168	158,000	+34,832
Research and Evaluation Prototypes	10,387	16,124	10,052	- 6,072
High Performance Network Facilities and Testbeds (ESnet)	28,293	29,722	30,000	+ 278
SBIR/STTR	—	6,194	6,857	+ 663
High Performance Computing and Network Facilities	*208,399*	*230,208*	*260,909*	*+30,701*
Total, Advanced Scientific Computing Research	**358,772***	**394,000**	**426,000**	**+32,000**

* Total reduced by $10,048,000: $8,972,000 of which was transferred to the Small Business Innovation Research (SBIR) program and $1,076,000 of which was transferred to the Small Business Technology Transfer (STTR) program.

ASCR Facilities

- Providing the Facility – High-End and Leadership Computing
- Investing in the Future - Research and Evaluation Prototypes
- Linking it all together – Energy Sciences Network (ESnet)

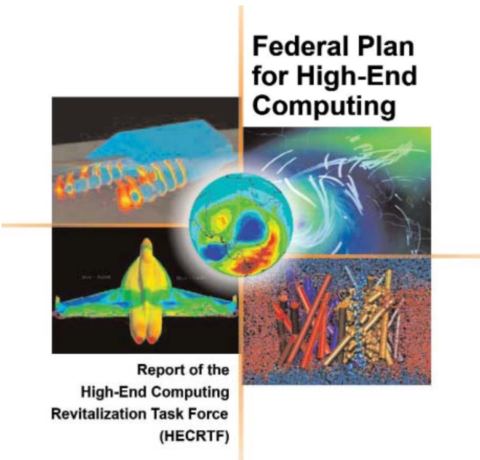

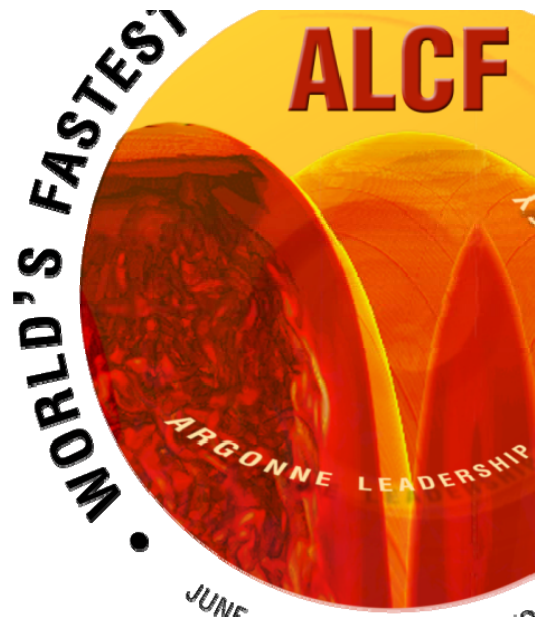

ASCR Research

Provide forefront research knowledge and foundational tools:

- Applied Mathematics
- Computer Science
- SciDAC
- Next Generation Networking for Science

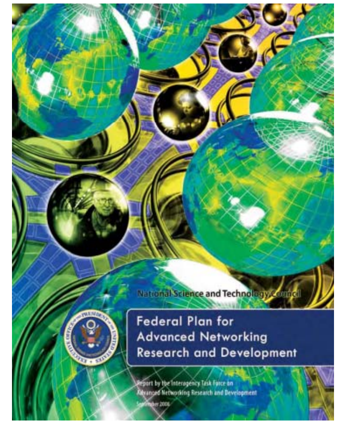

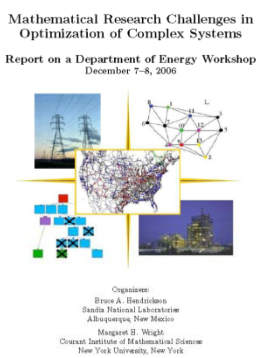

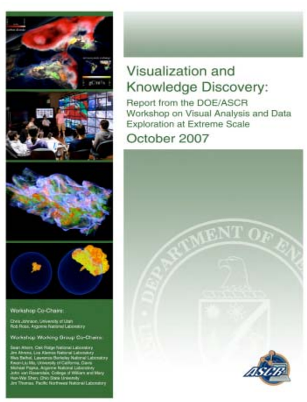

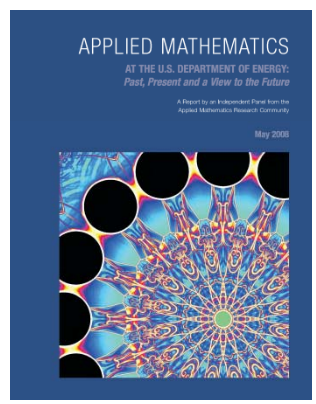

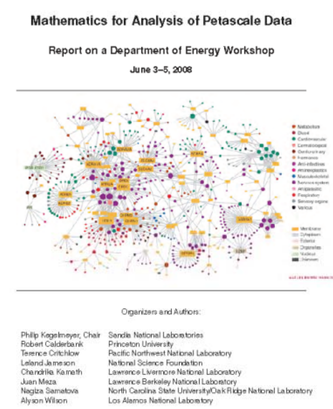

Applied Mathematics

(dollars in thousands)

	FY 2009	FY 2010	**FY 2011**	*FY 2011 vs FY 2010*	
Mathematical, Computational, and Computer Sciences Research					
Applied Mathematics	45,161	44,792	**45,450**	*+658*	***+1.5%***

- *What new mathematics are required to accurately model complex systems such as plasma containment in a Tokamak or the electric power grid taking place on vastly different time and/or length scales?*

Research Supported in 2010

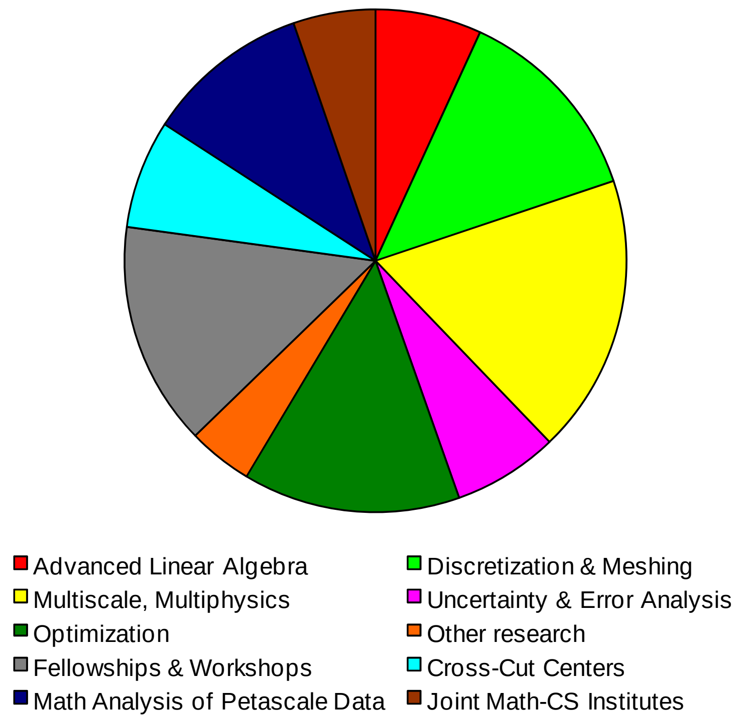

- **Numerical methods**
- **Advanced linear algebra**
- **Computational meshing**
- **Optimization of mathematical methods**
- **Multiscale mathematics and multiphysics**
- **Joint Applied Mathematics-Computer Science Institutes**
- **Mathematics for the analysis of extremely large datasets**
- **Mathematics of cyber security basic research**

- ***Computational Science Graduate Fellowship* and *Applied Mathematics and High Performance Computer Science Graduate Fellowship* programs.**

Computer Science

(dollars in thousands)

	FY 2009	FY 2010	**FY 2011**	*FY 2011 vs FY 2010*	
Mathematical, Computational, and Computer Sciences Research					
Computer Science	30,782	46,800	**47,400**	*+600*	***+1.3%***

- *What operating systems, data management, analysis, representation model development, user interface, and other tools are required to make effective use of current and future-generation supercomputers?*

Research supported in 2010

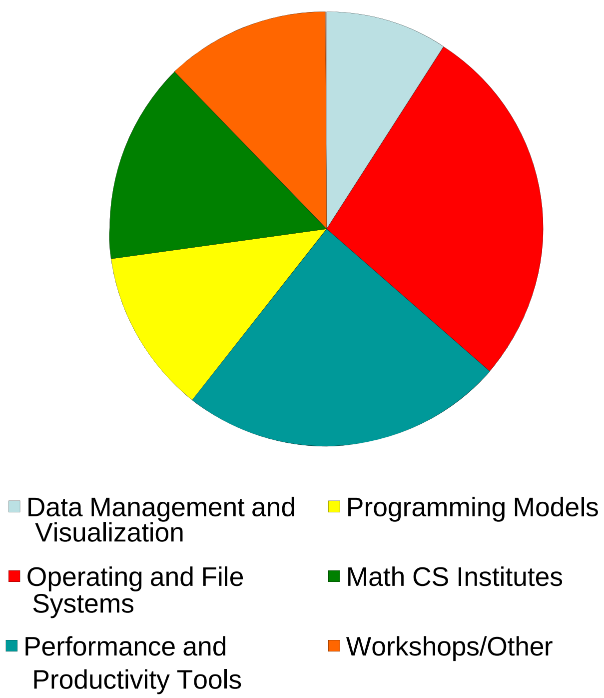

- **Operating and file systems.**
- **Performance and productivity tools**
- **Programming models**
- **Data management and visualization**
- **Joint Applied Mathematics-Computer Science Institutes**
- **Pre-competitive research in advanced computer architectures**
- **Applied Mathematics and High Performance Computer Science Graduate Fellowship programs**

Scientific Discovery through Advanced Computing

- Create comprehensive, scientific computing software infrastructure to enable scientific discovery in the physical, biological, and environmental sciences at the petascale

- Develop new generation of data management and knowledge discovery tools for large data sets (obtained from scientific user and simulations)

http://www.scidac.gov

U.S. DEPARTMENT OF ENERGY

Exascale Computing

Goal: Advance the Department's Science, Energy and National Security Missions through modeling and simulation at the extreme scale by the end of the decade

- **Provide up to 1,000x more powerful computing resources to**
  - **Advance scientific frontiers**
  - **Fully understand National & Global problems, their consequences, solutions and guide policy decisions**

- **Attributes**
  - **Integrated R&D project with software, hardware and application software**
  - **Broad community participation from universities, labs and industry such as computer vendors and chip manufactures**
  - **SC and NNSA partner**
  - **Coordinated with HPC efforts supported by NSA, DARPA, & NSF**

Exascale- Making the Case

- **Town Hall Meetings April-June 2007**
- **Scientific Grand Challenges Workshops Nov, 2008 – Oct, 2009**
  - Climate Science (11/08),
  - High Energy Physics (12/08),
  - Nuclear Physics (1/09),
  - Fusion Energy (3/09),
  - Nuclear Energy (5/09),
  - Biology (8/09),
  - Material Science and Chemistry (8/09),
  - National Security (10/09)
  - Cross-cutting technologies (2/10)
- **Exascale Steering Committee**
  - "Denver" vendor NDA visits 8/2009
  - SC09 vendor feedback meetings
  - Extreme Architecture and Technology Workshop 12/2009
- **International Exascale Software Project**

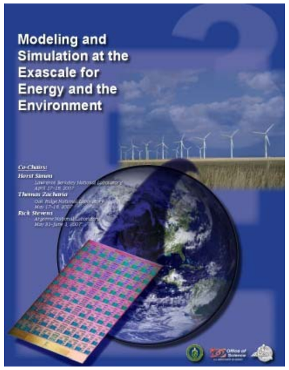

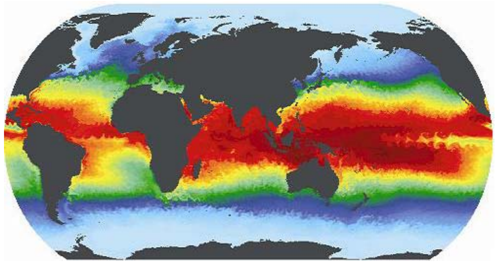

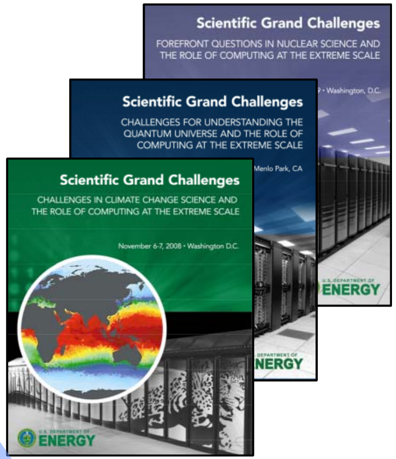

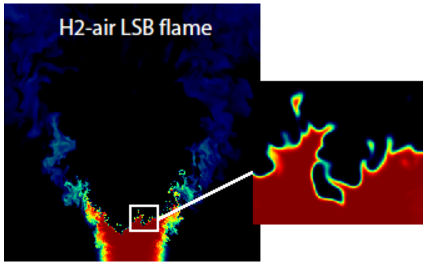

MISSION IMPERATIVES

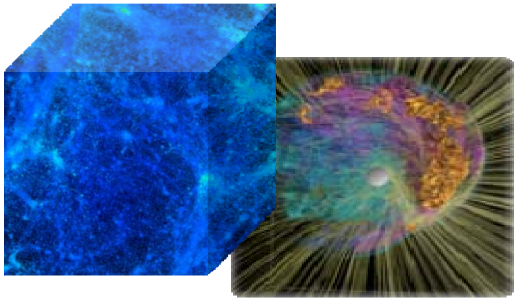

FUNDAMENTAL SCIENCE

Co- Design

Exascale simulation will enable fundamental advances in science

- High Energy & Nuclear Physics
  - Dark-energy and dark matter
  - Fundamentals of fission & fusion reactions
- Facility and experimental design
  - Effective design of accelerators
  - Probes of dark energy and dark matter
  - ITER shot planning and device control
- Materials / Chemistry
  - Predictive multi-scale materials modeling: observation to control
  - Effective, commercial technologies in renewable energy, catalysts, batteries and combustion
- Life Sciences
  - Better biofuels
  - Sequence to structure to function

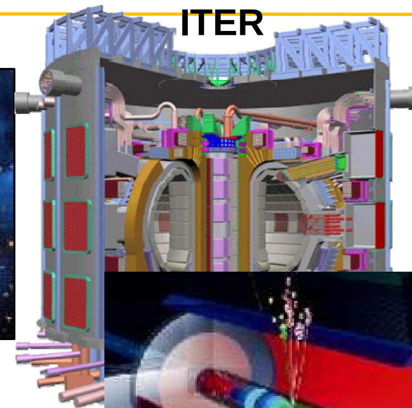

Hubble image of lensing

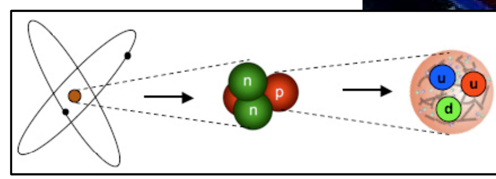

Structure of nucleons

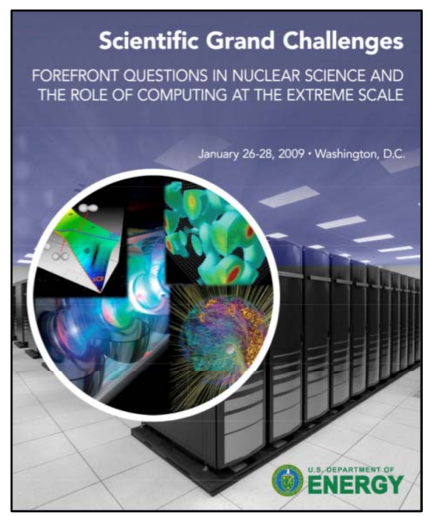

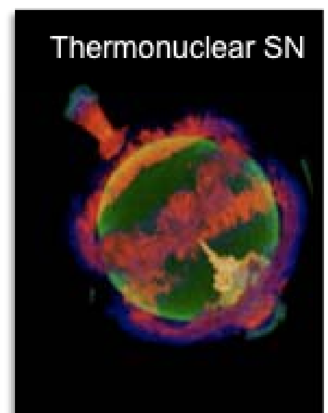

U.S. DEPARTMENT OF ENERGY

DOE Science and DOE Mission Rely on Cross-cutting Capabilities

- **Uncertainty quantification**
  - **Predict climate response to energy technology strategies**
  - **Assessment of safety, surety and performance of the aging/evolving stockpile without nuclear testing**
  - **Energy security**
  - **Responding to natural and manmade hazards**
- **Multi-scale, multi-physics modeling**
  - **Multiple physics packages in earth system model: ocean, land surface, atmosphere, ice**
  - **Multiple physics packages in modeling reactor core: neutronics, heat transfer, structures, fluids**
- **Statistics of rare events**
  - **Severe weather and surprises in climate system**
  - **Accident scenarios in nuclear energy**
  - **Nucleation of cracks and damage in materials**

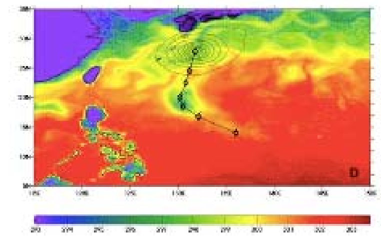
Storm track in ESM

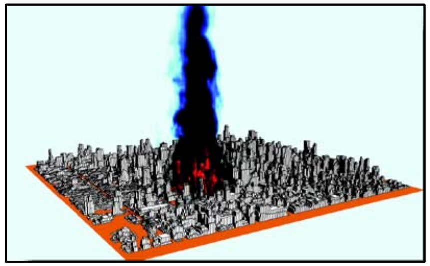
urban fire

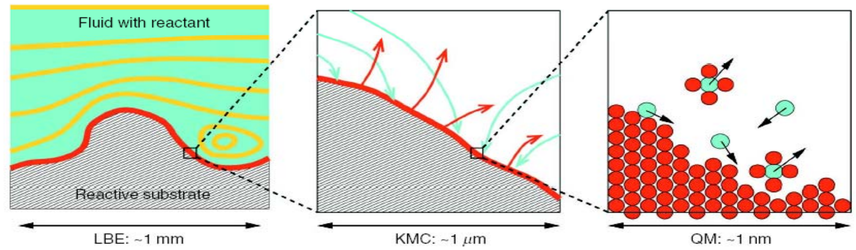

Multi-scale physics modeling

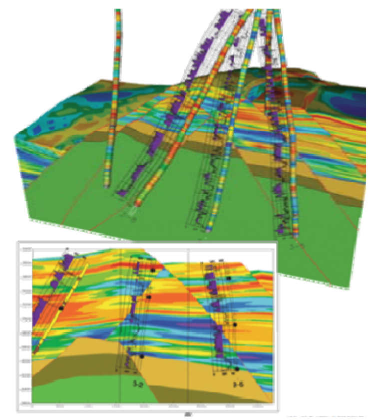
Geologic sequestration


Pine Island Glacier

Exascale ≠ Petascale X 1000

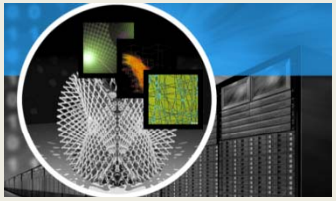

- **Traditionally, PDE-based applications have expected 10x increase in resolution with each 1000x increase in compute capability, *but don't expect it this time:***
  - We won't have 1000x the memory available
  - The processors won't be 10x faster
  - Proportionally, we won't be able to move as much data on or off each processor
  - Introduction of massive parallelism at the node level is a significant new challenge (MPI is only part of the solution)

- **However, exascale computing is an opportunity for…**
  - ***Scientific fidelity:*** Incorporate more physics instead of increased resolution
  - ***Greater Understanding:*** Develop uncertainty quantification (UQ) to establish confidence levels in computed results and deliver predictive science

U.S. DEPARTMENT OF ENERGY

Potential System Architecture

Systems	2009	2018	Difference Today & 2018
System peak	**2 Pflop/s**	**1 Eflop/s**	**O(1000)**
Power	**6 MW**	**~20 MW (goal)**	
System memory	0.3 PB	32 - 64 PB	O(100)
Node performance	125 GF	1,2 or 15TF	O(10) – O(100)
Node memory BW	25 GB/s	2 - 4TB/s	O(100)
Node concurrency	12	O(1k) or 10k	O(100) – O(1000)
Total Node Interconnect BW	3.5 GB/s	200-400GB/s (1:4 or 1:8 from memory BW)	O(100)
System size (nodes)	18,700	O(100,000) or O(1M)	O(10) – O(100)
Total concurrency	225,000	O(billion) [O(10) to O(100) for latency hiding]	O(10,000)
Storage	15 PB	500-1000 PB (>10x system memory is min)	O(10) – O(100)
IO	0.2 TB	60 TB/s	O(100)
MTTI	days	O(1 day)	- O(10)

Critical Challenges

- System power
- Memory
- Programming model
- Exascale processor designs
- Operating System strategy for exascale
- Reliability and resiliency

- HPC co-design strategy and implementation
  - Suite of a hierarchical performance models and simulators
  - Sustained commitment from applications, applied mathematics and computer science research, software developers and computer architects.

Positioning Applied Mathematics for Exascale

U.S. DEPARTMENT OF ENERGY

FY 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25

Drivers:

- Complex systems
- Massive data sets
- Emergence of many core, heterogeneous architectures
- Uncertainty quantification in large-scale applications
- Predictive Science
- Pervasive million-way concurrency architectures
- High-resolution simulations of energy systems
- Risk-informed decision-making through modeling and simulation

Applied Math Core Research:

- Multiscale math and numerical solutions of differential equations
- Scalable linear algebra for large systems
- Optimization algorithms and applications
- Meshing complex systems

New Applied Math Research Thrusts:

- Mathematics for Analysis of Petascale Data
- Joint Math/CS Institutes
- Complex Interconnected Systems
- Mathematics of Uncertainty Quantification
- Scalable modeling of complex, continuous and discrete systems
- Fault-tolerant , low memory bandwidth numerical methods
- Architecture-aware algorithms for multicomponent systems
- Mathematics for decision support through simulation

Application-Algorithm Challenges:

1PF 225K

10-20PF 3.2M

150PF ~50M

1-2EF ~1B

10EF

~10M-way concurrency

U.S. DEPARTMENT OF ENERGY

Computer Science Research: Approach for Creating Usable Exascale Environments for Computational Science

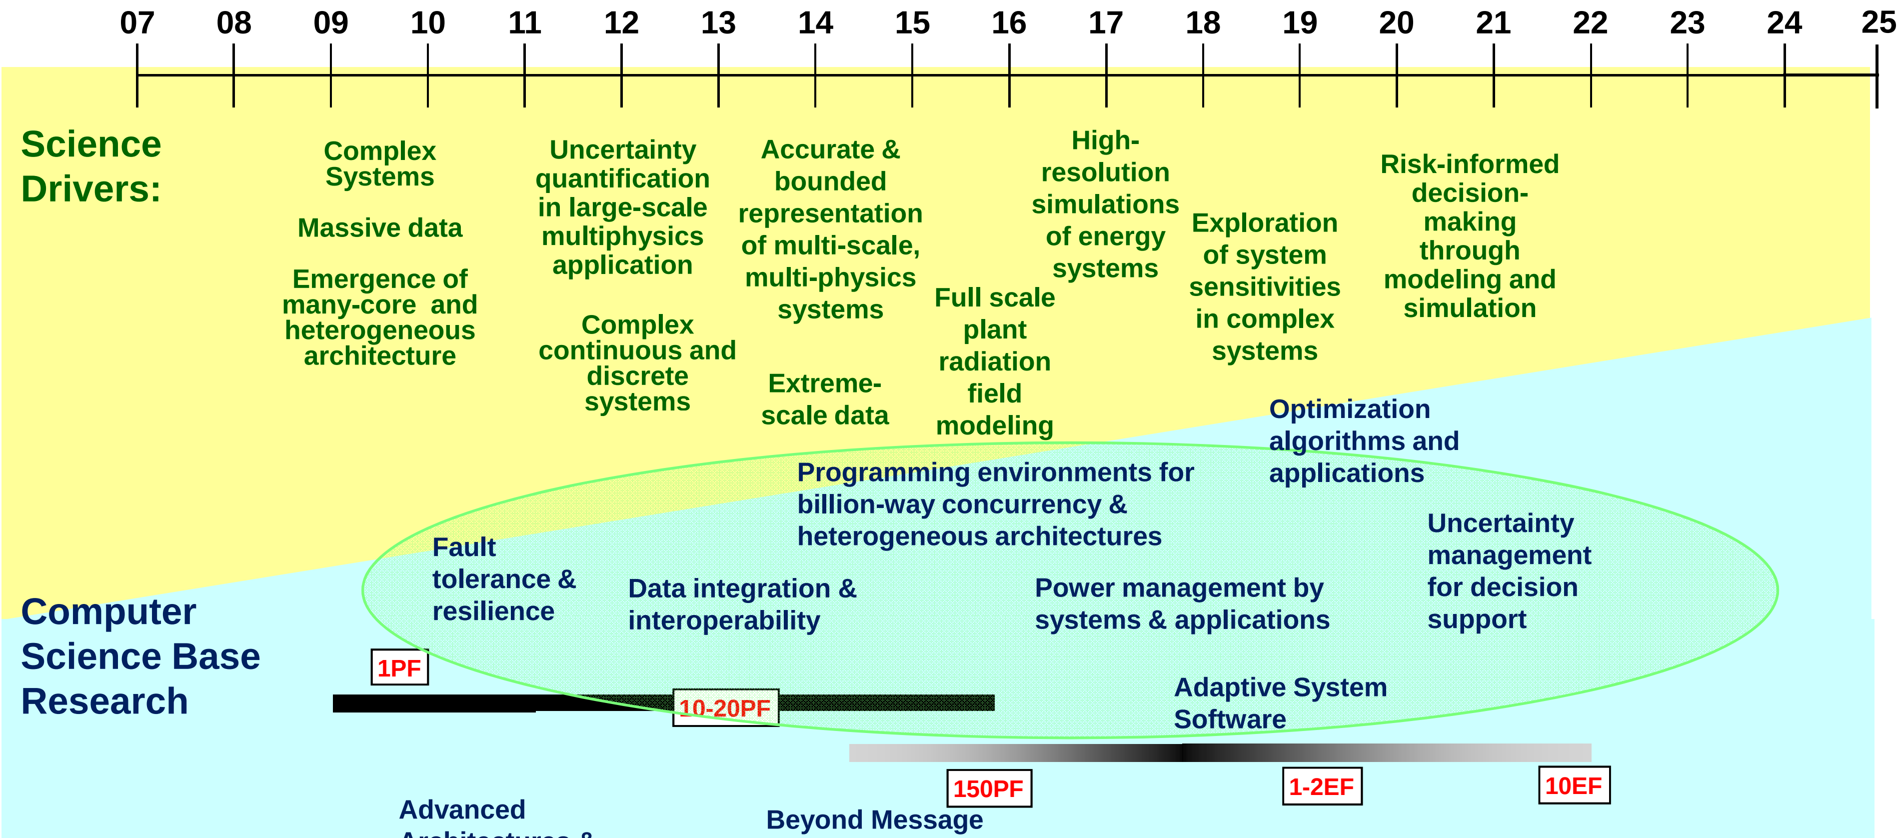

Computer Science Research Initiatives:

Joint Math/CS Institutes

Advanced Architectures & Technology

X-Stack

Data Management & Analysis for Extreme Scale Data

Assuring Simulation Credibility & Utility

Beyond Message Passing: Low-Power Synchronization

Paradigms for Scalability

I/O Management, Data Organization & Integration

Scientist-in-the-loop interactive simulation steering & data analysis

Autotuning & Fault Management

Emergent Area Research

U.S. DEPARTMENT OF ENERGY

Summary

- Extreme Scale systems projected for 2015 – 2020 will need fundamental changes in Execution Model and System Software to address Concurrency, Energy Efficiency, and Resiliency challenges
- Applications must exploit a combination of strong scaling, traditional weak scaling, and new-era weak scaling techniques, in conjunction with suitable attention to energy efficiency (data movement & locality)
- Programming models will need to express all intrinsic parallelism and locality for forward scalability and portability
- System software will need to be co-designed across multiple levels and with hardware for effective management of locality and parallelism in Extreme Scale systems
    - Current & future many-core systems can serve as useful intermediate platforms

Regarding Exascale…

… maybe Mark Twain has it right

"Twenty years from now you will be more disappointed by the things that you didn't do than by the ones you did do. So throw off the bowlines. Sail away from the safe harbor. Catch the trade winds in your sails. Explore. Dream. Discover."